

# Optimal detection of changepoints with a linear computational cost

Killick, R., Fearnhead, P. and Eckley, I.A.\*

October 10, 2012

## Abstract

We consider the problem of detecting multiple changepoints in large data sets. Our focus is on applications where the number of changepoints will increase as we collect more data: for example in genetics as we analyse larger regions of the genome, or in finance as we observe time-series over longer periods. We consider the common approach of detecting changepoints through minimising a cost function over possible numbers and locations of changepoints. This includes several established procedures for detecting changing points, such as penalised likelihood and minimum description length. We introduce a new

---

\*R. Killick is Senior Research Associate, Department of Mathematics & Statistics, Lancaster University, Lancaster, UK (E-mail: r.killick@lancs.ac.uk). P. Fearnhead is Professor, Department of Mathematics & Statistics, Lancaster University, Lancaster, UK (E-mail: p.fearnhead@lancs.ac.uk). I.A. Eckley is Senior Lecturer, Department of Mathematics & Statistics, Lancaster University, Lancaster, UK (E-mail: i.eckley@lancs.ac.uk). The authors are grateful to Richard Davis and Alice Cleynen for providing the Auto-PARM and PDPA software respectively. Part of this research was conducted whilst R. Killick was a jointly funded Engineering and Physical Sciences Research Council (EPSRC) / Shell Research Ltd graduate student at Lancaster University. Both I.A. Eckley and R. Killick also gratefully acknowledge the financial support of the EPSRC grant number EP/I016368/1.

method for finding the minimum of such cost functions and hence the optimal number and location of changepoints that has a computational cost which, under mild conditions, is linear in the number of observations. This compares favourably with existing methods for the same problem whose computational cost can be quadratic or even cubic. In simulation studies we show that our new method can be orders of magnitude faster than these alternative exact methods. We also compare with the Binary Segmentation algorithm for identifying changepoints, showing that the exactness of our approach can lead to substantial improvements in the accuracy of the inferred segmentation of the data.

**KEYWORDS** Structural Change; Dynamic Programming; Segmentation; PELT.

## 1. INTRODUCTION

As increasingly longer data sets are being collected, more and more applications require the detection of changes in the distributional properties of such data. Consider for example recent work in genomics, looking at detecting changes in gene copy numbers or in the compositional structure of the genome (Braun et al., 2000; Olshen et al., 2004; Picard et al., 2005); and in finance where, for example, interest lies in detecting changes in the volatility of time series (Aggarwal et al., 1999; Andreou and Ghysels, 2002; Fernandez, 2004). Typically such series will contain several changepoints. There is therefore a growing need to be able to search for such changes efficiently. It is this search problem which we consider in this paper. In particular we focus on applications where we expect the number of changepoints to increase as we collect more data. This is a natural assumption in many cases, for example as we analyse longer regions of the genome or as we record financial time-series over longer time-periods. By comparison it does not necessarily apply to situations where we are obtaining data over a fixed time-period at a higher frequency.

At the time of writing Binary Segmentation proposed by Scott and Knott (1974) is arguably the most widely used changepoint search method. It is approximate in nature with an  $\mathcal{O}(n \log n)$  computational cost, where  $n$  is the number of data points. While exact search algorithms exist for the most common forms of changepoint models, these have a much greater computational cost. Several exact search methods are based on dynamic programming. For example the Segment Neighbourhood method proposed by Auger and Lawrence (1989) is  $\mathcal{O}(Qn^2)$ , where  $Q$  is the maximum number of changepoints you wish to search for. Note that in scenarios where the number of changepoints increases linearly with  $n$ , this can correspond to a computational cost that is cubic in the length of the data. An alternative dynamic programming algorithm is provided by the Optimal Partitioning approach of Jackson et al. (2005). As we describe in Section 2.2 this can be applied to a slightly smaller class of problems and is an exact approach whose computational cost is  $\mathcal{O}(n^2)$ .

We present a new approach to search for changepoints, which is exact and under mild conditions has a computational cost that is linear in the number of data points: the **Pruned Exact Linear Time** (PELT) method. This approach is based on the algorithm of Jackson et al. (2005), but involves a pruning step within the dynamic program. This pruning reduces the computational cost of the method, but does not affect the exactness of the resulting segmentation. It can be applied to find changepoints under a range of statistical criteria such as penalised likelihood, quasi-likelihood (Braun et al., 2000) and cumulative sum of squares (Inclan and Tiao, 1994; Picard et al., 2011). In simulations we compare PELT with both Binary Segmentation and Optimal Partitioning. We show that PELT can be calculated orders of magnitude faster than Optimal Partitioning, particularly for long data sets. Whilst asymptotically PELT can be quicker, we find that in practice Binary Segmentation is quicker on the examples we consider, and we believe this would be the case in almost all applications.

However, we show that PELT leads to a substantially more accurate segmentation than Binary Segmentation.

The paper is organised as follows. We begin in Section 2 by reviewing some basic changepoint notation and summarizing existing work in the area of search methods. The PELT method is introduced in Section 3 and the computational cost of this approach is considered in Section 3.1. The efficiency and accuracy of the PELT method are demonstrated in Section 4. In particular we demonstrate the methods' performance on large data sets coming from oceanographic (Section 4.2) and financial (supplementary material) applications. Results show the speed gains over other exact search methods and the increased accuracy relative to approximate search methods such as Binary Segmentation. The paper concludes with a discussion.

## 2. BACKGROUND

Changepoint analysis can, loosely speaking, be considered to be the identification of points within a data set where the statistical properties change. More formally, let us assume we have an ordered sequence of data,  $y_{1:n} = (y_1, \dots, y_n)$ . Our model will have a number of changepoints,  $m$ , together with their positions,  $\tau_{1:m} = (\tau_1, \dots, \tau_m)$ . Each changepoint position is an integer between 1 and  $n - 1$  inclusive. We define  $\tau_0 = 0$  and  $\tau_{m+1} = n$  and assume that the changepoints are ordered such that  $\tau_i < \tau_j$  if, and only if,  $i < j$ . Consequently the  $m$  changepoints will split the data into  $m + 1$  segments, with the  $i$ th segment containing  $y_{(\tau_{i-1}+1):\tau_i}$ .

One commonly used approach to identify multiple changepoints is to minimise:

$$\sum_{i=1}^{m+1} [\mathcal{C}(y_{(\tau_{i-1}+1):\tau_i})] + \beta f(m). \quad (1)$$

Here  $\mathcal{C}$  is a cost function for a segment and  $\beta f(m)$  is a penalty to guard against over fitting. Twice the negative log likelihood is a commonly used cost function in the changepoint literature (see for example Horvath, 1993; Chen and Gupta, 2000),

although other cost functions such as quadratic loss and cumulative sums are also used (e.g. Rigaiill, 2010; Inlan and Tiao, 1994), or those based on both the segment log-likelihood and the length of the segment (Zhang and Siegmund, 2007). Turning to choice of penalty, in practice by far the most common choice is one which is linear in the number of changepoints, i.e.  $\beta f(m) = \beta m$ . Examples of such penalties include Akaike’s Information Criterion (AIC, Akaike (1974)) ( $\beta = 2p$ ) and Schwarz Information Criterion (SIC, also known as BIC; Schwarz, 1978) ( $\beta = p \log n$ ), where  $p$  is the number of additional parameters introduced by adding a changepoint. The PELT method which we introduce in Section 3 is designed for such linear cost functions. Although linear cost functions are commonplace within the changepoint literature Guyon and Yao (1999), Picard et al. (2005) and Birge and Massart (2007) offer examples and discussion of alternative penalty choices. In Section 3.2 we show how PELT can be applied to some of these alternative choices.

The remainder of this section describes two commonly used methods for multiple changepoint detection; Binary Segmentation (Scott and Knott, 1974) and Segment Neighbourhoods (Auger and Lawrence, 1989). A third method proposed by Jackson et al. (2005) is also described as it forms the basis for the PELT method which we propose. For notational simplicity we describe all the algorithms (including PELT) assuming that the minimum segment length is a single observation, i.e.  $\tau_{i-1} - \tau_i \geq 1$ . A larger minimum segment length is easily implemented when appropriate, see for example Section 4.

## 2.1 Binary Segmentation

Binary Segmentation (BS) is arguably the most established search method used within the changepoint literature. Early applications include Scott and Knott (1974) and Sen and Srivastava (1975). In essence the method extends any single changepoint method

to multiple changepoints by iteratively repeating the method on different subsets of the sequence. It begins by initially applying the single changepoint method to the entire data set, i.e. we test if a  $\tau$  exists that satisfies

$$\mathcal{C}(y_{1:\tau}) + \mathcal{C}(y_{(\tau+1):n}) + \beta < \mathcal{C}(y_{1:n}). \quad (2)$$

If (2) is false then no changepoint is detected and the method stops. Otherwise the data is split into two segments consisting of the sequence before and after the identified changepoint,  $\tau_a$  say, and apply the detection method to each new segment. If either or both tests are true, we split these into further segments at the newly identified changepoint(s), applying the detection method to each new segment. This procedure is repeated until no further changepoints are detected. For pseudo-code of the BS method see for example Eckley et al. (2011).

Binary Segmentation can be viewed as attempting to minimise equation (1) with  $f(m) = m$ : each step of the algorithm attempts to introduce an extra changepoint if and only if it reduces (1). The advantage of the BS method is that it is computationally efficient, resulting in an  $\mathcal{O}(n \log n)$  calculation. However this comes at a cost as it is not guaranteed to find the global minimum of (1).

## 2.2 Exact methods

***Segment Neighbourhood*** Auger and Lawrence (1989) propose an alternative, exact search method for changepoint detection, namely the Segment Neighbourhood (SN) method. This approach searches the entire segmentation space using dynamic programming (Bellman and Dreyfus, 1962). It begins by setting an upper limit on the size of the segmentation space (i.e. the maximum number of changepoints) that is required – this is denoted  $Q$ . The method then continues by computing the cost function for all possible segments. From this all possible segmentations with between 0 and  $Q$  changepoints are considered.

In addition to being an exact search method, the SN approach has the ability to incorporate an arbitrary penalty of the form,  $\beta f(m)$ . However, a consequence of the exhaustive search is that the method has significant computational cost,  $\mathcal{O}(Qn^2)$ . If as the observed data increases, the number of changepoints increases linearly, then  $Q = \mathcal{O}(n)$  and the method will have a computational cost of  $\mathcal{O}(n^3)$ .

**The optimal partitioning method** Yao (1984) and Jackson et al. (2005) propose a search method that aims to minimise

$$\sum_{i=1}^{m+1} [\mathcal{C}(y_{(\tau_{i-1}+1):\tau_i}) + \beta]. \quad (3)$$

This is equivalent to (1) where  $f(m) = m$ .

Following Jackson et al. (2005) the optimal partitioning (OP) method begins by first conditioning on the last point of change. It then relates the optimal value of the cost function to the cost for the optimal partition of the data prior to the last changepoint plus the cost for the segment from the last changepoint to the end of the data. More formally, let  $F(s)$  denote the minimisation from (3) for data  $y_{1:s}$  and  $\mathcal{T}_s = \{\tau : 0 = \tau_0 < \tau_1 < \dots < \tau_m < \tau_{m+1} = s\}$  be the set of possible vectors of changepoints for such data. Finally set  $F(0) = -\beta$ . It therefore follows that:

$$\begin{aligned} F(s) &= \min_{\tau \in \mathcal{T}_s} \left\{ \sum_{i=1}^{m+1} [\mathcal{C}(y_{(\tau_{i-1}+1):\tau_i}) + \beta] \right\}, \\ &= \min_t \left\{ \min_{\tau \in \mathcal{T}_t} \sum_{i=1}^m [\mathcal{C}(y_{(\tau_{i-1}+1):\tau_i}) + \beta] + \mathcal{C}(y_{(t+1):n}) + \beta \right\}, \\ &= \min_t \{ F(t) + \mathcal{C}(y_{(t+1):n}) + \beta \}. \end{aligned}$$

This provides a recursion which gives the minimal cost for data  $y_{1:s}$  in terms of the minimal cost for data  $y_{1:t}$  for  $t < s$ . This recursion can be solved in turn for  $s = 1, 2, \dots, n$ . The cost of solving the recursion for time  $s$  is linear in  $s$ , so the overall computational cost of finding  $F(n)$  is quadratic in  $n$ . Steps for implementing the OP method are given in Algorithm 1.

---

### Optimal Partitioning

**Input:** A set of data of the form,  $(y_1, y_2, \dots, y_n)$  where  $y_i \in \mathbb{R}$ .

A measure of fit  $\mathcal{C}(\cdot)$  dependent on the data.

A penalty constant  $\beta$  which does not depend on the number or location of changepoints.

**Initialise:** Let  $n$  = length of data and set  $F(0) = -\beta$ ,  $cp(0) = NULL$ .

**Iterate** for  $\tau^* = 1, \dots, n$

1. Calculate  $F(\tau^*) = \min_{0 \leq \tau < \tau^*} [F(\tau) + \mathcal{C}(y_{(\tau+1):\tau^*}) + \beta]$ .
2. Let  $\tau' = \arg \{ \min_{0 \leq \tau < \tau^*} [F(\tau) + \mathcal{C}(y_{(\tau+1):\tau^*}) + \beta] \}$ .
3. Set  $cp(\tau^*) = (cp(\tau'), \tau')$ .

**Output** the change points recorded in  $cp(n)$ .

---

### Algorithm 1: Optimal Partitioning.

Whilst OP improves on the computational efficiency of the SN method, it is still far from being competitive computationally with the BS method. Section 3 introduces a modification of the optimal partitioning method denoted PELT which results in an approach whose computational cost can be linear in  $n$  whilst retaining an exact minimisation of (3). This exact and efficient computation is achieved via a combination of optimal partitioning and pruning.

## 3. A PRUNED EXACT LINEAR TIME METHOD

We now consider how pruning can be used to increase the computational efficiency of the OP method whilst still ensuring that the method finds a global minimum of the cost function (3). The essence of pruning in this context is to remove those values of  $\tau$  which can never be minima from the minimisation performed at each iteration in (1) of Algorithm 1.

The following theorem gives a simple condition under which we can do such pruning.

**Theorem 3.1** *We assume that when introducing a changepoint into a sequence of observations the cost,  $\mathcal{C}$ , of the sequence reduces. More formally, we assume there exists a constant  $K$  such that for all  $t < s < T$ ,*

$$\mathcal{C}(y_{(t+1):s}) + \mathcal{C}(y_{(s+1):T}) + K \leq \mathcal{C}(y_{(t+1):T}). \quad (4)$$

*Then if*

$$F(t) + \mathcal{C}(y_{(t+1):s}) + K \geq F(s) \quad (5)$$

*holds, at a future time  $T > s$ ,  $t$  can never be the optimal last changepoint prior to  $T$ .*

**Proof.** See Section 5 of Supplementary Material.  $\square$

The intuition behind this result is that if (5) holds then for any  $T > s$  the best segmentation with the most recent changepoint prior to  $T$  being at  $s$  will be better than any which has this most recent changepoint at  $t$ . Note that almost all cost functions used in practice satisfy assumption (4). For example, if we take the cost function to be minus the log-likelihood then the constant  $K = 0$  and if we take it to be minus a penalised log-likelihood then  $K$  would equal the penalisation factor.

The condition imposed in Theorem 3.1 for a candidate changepoint,  $t$ , to be discarded from future consideration is important as it removes computations that are not relevant for obtaining the final set of changepoints. This condition can be easily implemented into the OP method and the pseudo-code is given in Algorithm 2. This shows that at each step in the method the candidate changepoints satisfying the condition are noted and removed from the next iteration. We show in the next section that under certain conditions the computational cost of this method will be linear in the number of observations, as a result we call this the **Pruned Exact Linear Time** (PELT) method.

---

### PELT Method

**Input:** A set of data of the form,  $(y_1, y_2, \dots, y_n)$  where  $y_i \in \mathbb{R}$ .

A measure of fit  $\mathcal{C}(\cdot)$  dependent on the data.

A penalty constant  $\beta$  which does not depend on the number or location of changepoints.

A constant  $K$  that satisfies equation 4.

**Initialise:** Let  $n$  = length of data and set  $F(0) = -\beta$ ,  $cp(0) = NULL$ ,  $R_1 = \{0\}$ .

**Iterate** for  $\tau^* = 1, \dots, n$

1. Calculate  $F(\tau^*) = \min_{\tau \in R_{\tau^*}} [F(\tau) + \mathcal{C}(y_{(\tau+1):\tau^*}) + \beta]$ .
2. Let  $\tau^1 = \arg \{ \min_{\tau \in R_{\tau^*}} [F(\tau) + \mathcal{C}(y_{(\tau+1):\tau^*}) + \beta] \}$ .
3. Set  $cp(\tau^*) = [cp(\tau^1), \tau^1]$ .
4. Set  $R_{\tau^*+1} = \{ \tau \in R_{\tau^*} \cup \{\tau^*\} : F(\tau) + \mathcal{C}(y_{\tau+1:\tau^*}) + K \leq F(\tau^*) \}$ .

**Output** the change points recorded in  $cp(n)$ .

---

### Algorithm 2: PELT Method.

#### 3.1 Linear Computational Cost of PELT

We now investigate the theoretical computational cost of the PELT method. We focus on the most important class of changepoint models and penalties and provide sufficient conditions for the method to have a computational cost that is linear in the number of data points. The case we focus on is the set of models where the segment parameters are independent across segments and the cost function for a segment is minus the maximum log-likelihood value for the data in that segment.

More formally, our result relates to the expected computational cost of the method and how this depends on the number of data points we analyse. To this end we define an underlying stochastic model for the data generating process. Specifically we define such a process over positive-integer time points and then consider analysing the first  $n$  data points generated by this process. Our result assumes that the parameters

associated with a given segment are IID with density function  $\pi(\theta)$ . For notational simplicity we assume that given the parameter,  $\theta$ , for a segment, the data points within the segment are IID with density function  $f(y|\theta)$  although extensions to dependence within a segment is trivial. Finally, as previously stated our cost function will be based on minus the maximum log-likelihood:

$$\mathcal{C}(y_{(t+1):s}) = -\max_{\theta} \sum_{i=t+1}^s \log f(y_i|\theta).$$

Note that for this loss-function,  $K = 0$  in (4). Hence pruning in PELT will just depend on the choice of penalty constant  $\beta$ .

We also require a stochastic model for the location of the changepoints in the form of a model for the length of each segment. If the changepoint positions are  $\tau_1, \tau_2, \dots$ , then define the segment lengths to be  $S_1 = \tau_1$  and for  $i = 2, 3, \dots$ ,  $S_i = \tau_i - \tau_{i-1}$ . We assume the  $S_i$  are IID copies of a random variable  $S$ . Furthermore  $S_1, S_2, \dots$ , are independent of the parameters associated with the segments.

**Theorem 3.2** *Define  $\theta^*$  to be the value that maximises the expected log-likelihood*

$$\theta^* = \arg \max \int \int f(y|\theta) f(y|\theta_0) dy \pi(\theta_0) d\theta_0.$$

*Let  $\theta_i$  be the true parameter associated with the segment containing  $y_i$  and  $\hat{\theta}_n$  be the maximum likelihood estimate for  $\theta$  given data  $y_{1:n}$  and an assumption of a single segment:*

$$\hat{\theta}_n = \arg \max_{\theta} \sum_{i=1}^n \log f(y_i|\theta).$$

*Then if*

*(A1) denoting*

$$B_n = \sum_{i=1}^n \left[ \log f(y_i|\hat{\theta}_n) - \log f(y_i|\theta^*) \right],$$

*we have  $\mathbb{E}(B_n) = o(n)$  and  $\mathbb{E}([B_n - \mathbb{E}(B_n)]^4) = \mathcal{O}(n^2)$ ;*

(A2)

$$\mathbb{E} \left( [\log f(Y_i|\theta_i) - \log f(Y_i|\theta^*)]^4 \right) < \infty;$$

(A3)

$$\mathbb{E} (S^4) < \infty; \text{ and}$$

(A4)

$$\mathbb{E} (\log f(Y_i|\theta_i) - \log f(Y_i|\theta^*)) > \frac{\beta}{\mathbb{E}(S)};$$

where  $S$  is the expected segment length, the expected CPU cost of PELT for analysing  $n$  data points is bounded above by  $Ln$  for some constant  $L < \infty$ .

**Proof.** See Section 6 of Supplementary Material. □

Conditions (A1) and (A2) of Theorem 3.2 are weak technical conditions. For example, general asymptotic results for maximum likelihood estimation would give  $B_n = \mathcal{O}_p(1)$ , and (A1) is a slightly stronger condition which is controlling the probability of  $B_n$  taking values that are  $\mathcal{O}(n^{1/2})$  or greater.

The other two conditions are more important. Condition (A3) is needed to control the probability of large segments. One important consequence of (A3) is that the expected number of changepoints will increase linearly with  $n$ . Finally condition (A4) is a natural one as it is required for the expected penalised likelihood value obtained with the true changepoint and parameter values to be greater than the expected penalised likelihood value if we fit a single segment to the data with segment parameter  $\theta^*$ .

In all cases the worst case complexity of the algorithm is where no pruning occurs and the computational cost is  $\mathcal{O}(n^2)$ .

### 3.2 PELT for concave penalties

There is a growing body of research (see Guyon and Yao, 1999; Picard et al., 2005; Birge and Massart, 2007) that consider nonlinear penalty forms. In this section we

address how PELT can be applied to penalty functions which are concave.

$$\beta f(m) + \sum_{i=1}^{m+1} \mathcal{C}(y_{(\tau_{i-1}+1):\tau_i}), \quad (6)$$

where  $f(m)$  is concave and differentiable.

For an appropriately chosen  $\gamma$ , the following result shows that the optimum segmentation based on such a penalty corresponds to minimising

$$m\gamma + \sum_{i=1}^{m+1} \mathcal{C}(y_{(\tau_{i-1}+1):\tau_i}). \quad (7)$$

**Theorem 3.3** *Assume that  $f$  is concave and differentiable, with derivative denoted  $f'$ . Further, let  $\hat{m}$  be the value of  $m$  for which the criteria (6) is minimised. Then the optimal segmentation under this set of penalties is the segmentation that minimises*

$$mf'(\hat{m}) + \sum_{i=1}^{m+1} \mathcal{C}(y_{(\tau_{i-1}+1):\tau_i}). \quad (8)$$

**Proof.** See Section 7 of Supplementary Material.  $\square$

This suggests that we can minimize penalty functions based on  $f(m)$  using PELT – the correct penalty constant just needs to be applied. A simple approach, is to run PELT with an arbitrary penalty constant, say  $\gamma = f'(1)$ . Let  $m_0$  denote the resulting number of changepoints estimated. We then run PELT with penalty constant  $\gamma = f'(m_0)$ , and get a new estimate of the number of changepoints  $m_1$ . If  $m_0 = m_1$  we stop. Otherwise we update the penalty constant and repeat until convergence. This simple procedure is not guaranteed to find the optimal number of changepoints. Indeed more elaborate search schemes may be better. However, as tests of this simple approach in Section 4.3 show, it can be quite effective.

#### 4. SIMULATION AND DATA EXAMPLES

We now compare PELT with both Optimal Partitioning (OP) and Binary Segmentation (BS) on a range of simulated and real examples. Our aim is to see empirically

(i) how the computational cost of PELT is affected by the amount of data, (ii) to evaluate the computational savings that PELT gives over OP, and (iii) to evaluate the increased accuracy of exact methods over BS. Unless otherwise stated, we used the SIC penalty. In this case the penalty constant increases with the amount of data, and as such the application of PELT lies outside the conditions of Theorem 3.2. We also consider the impact of the number of changepoints not increasing linearly with the amount of data, a further violation of the conditions of Theorem 3.2.

#### 4.1 Changes in Variance within Normally Distributed Data

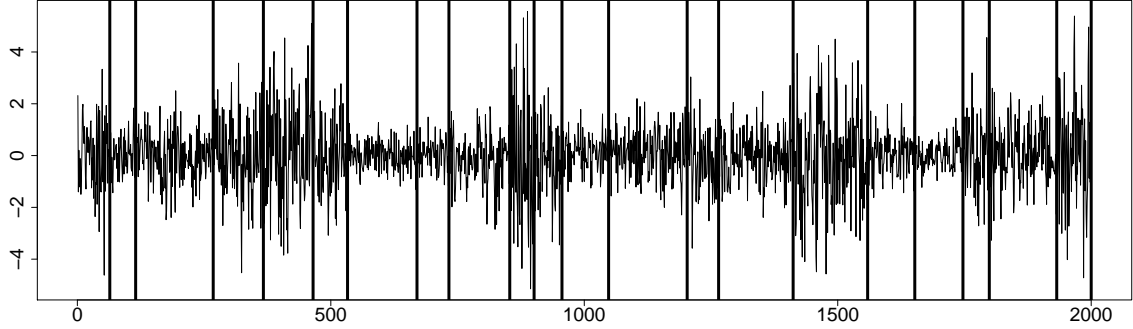
In the following subsections we consider multiple changes in variance within data sets that are assumed to follow a Normal distribution with a constant (unknown) mean. We begin by showing the power of the PELT method in detecting multiple changes via a simulation study, and then use PELT to analyse Oceanographic data and Dow Jones Index returns (Section 2 in supplementary material).

***Simulation Study*** In order to evaluate PELT we shall construct sets of simulated data on which we shall run various multiple changepoint methods. It is reasonable to set the cost function,  $\mathcal{C}$  as twice the negative log-likelihood. Note that for a change in variance (with unknown mean), the minimum segment length is two observations. The cost of a segment is then

$$\mathcal{C}(y_{(\tau_{i-1}+1):\tau_i}) = (\tau_i - \tau_{i-1}) \left( \log(2\pi) + \log \left( \frac{\sum_{j=\tau_{i-1}+1}^{\tau_i} (y_j - \mu)^2}{\tau_i - \tau_{i-1}} \right) + 1 \right). \quad (9)$$

Our simulated data consists of scenarios with varying lengths,  $n = (100, 200, 500, 1000, 2000, 5000, 10000, 20000, 50000)$ . For each value of  $n$  we consider a linearly increasing number of changepoints,  $m = n/100$ . In each case the changepoints are distributed uniformly across  $(2, n-2)$  with the only constraint being that there must be at least 30 observations between changepoints. Within each of these scenarios

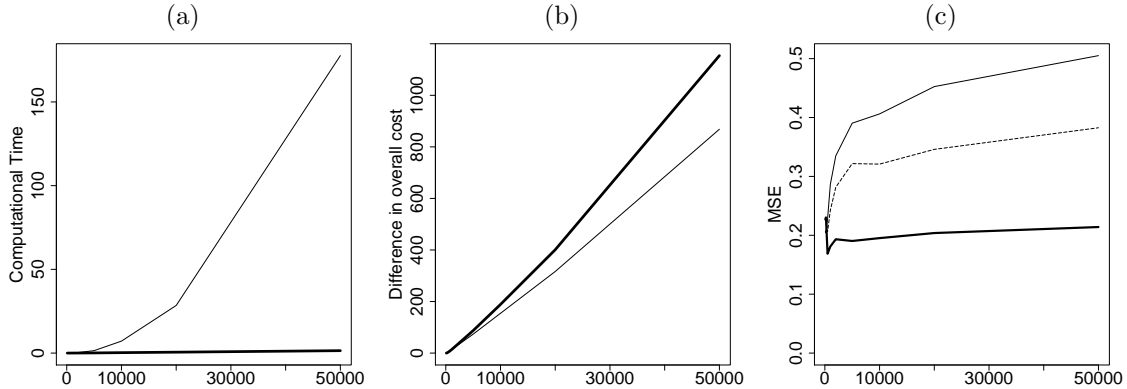
Figure 1: A realisation of multiple changes in variance where the true changepoint locations are shown by vertical lines.



we have 1,000 repetitions where the mean is fixed at 0 and the variance parameters for each segment are assumed to have a Log-Normal distribution with mean 0 and standard deviation  $\frac{\log(10)}{2}$ . These parameters are chosen so that 95% of the simulated variances are within the range  $[\frac{1}{10}, 10]$ . An example realisation is shown in Figure 1. Additional simulations considering a wider range of options for the number of changepoints (square root:  $m = \lfloor \sqrt{n}/4 \rfloor$  and fixed:  $m = 2$ ) and parameter values are given in Section 1 of the Supplementary Material.

Results are shown in Figure 2 where we denote the Binary Segmentation method which identifies the *same* number of changepoints as PELT as subBS. Conversely the number of changepoints BS would optimally select is called optimal BS. Firstly Figure 2(a) shows that when the number of changepoints increases linearly with  $n$ , PELT does indeed have a CPU cost that is linear in  $n$ . By comparison figures in the supplementary material show that if the number of changepoints increases at a slower rate, for example, square root or even fixed number of changepoints, the CPU cost of PELT is no longer linear. However even in the latter two cases, substantial computational savings are attained relative to OP. Comparison of times with BS are also given in the Supplementary material. These show that PELT and BS have similar computational costs for the case of linearly increasing number of changepoints, but

Figure 2: (a) Average Computational Time (in seconds) for a change in variance (thin: OP, thick: PELT). (b) Average difference in cost between PELT and BS for subBS (thin), optimal BS (thick). (c) MSE for PELT (thick), optimal BS (thin) and subBS (dotted).



BS can be orders of magnitude quicker for other situations.

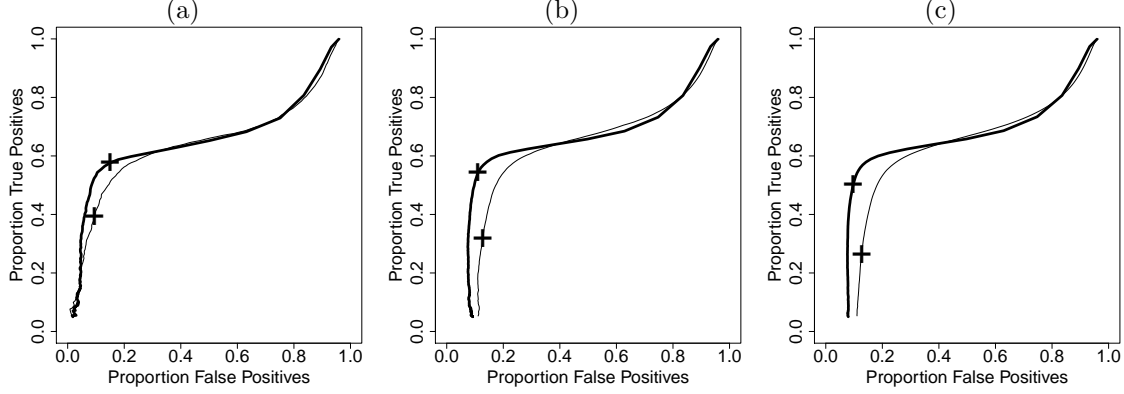
The advantage of PELT over BS is that PELT is guaranteed to find the optimal segmentation under the chosen cost function, and as such is likely to be preferred providing sufficient computational time is available to run it. Figure 2(b) shows the improved fit to the data that PELT attains over BS in terms of the smaller values of the cost function that are found. If you consider using the log-likelihood to choose between competing models, the value for  $n = 50,000$  is over 1000 which is very large. An alternative comparison is to look at how well each method estimates the parameters in the model. We measure this using mean square error:

$$\text{MSE} = \frac{\sum_{i=1}^n (\hat{\theta}_i - \theta_i)^2}{n}, \quad (10)$$

Figure 2(c) shows the increase in accuracy in terms of mean square error of estimates of the parameter. The figures in the supplementary material show that for the fixed number of changepoints scenario the difference is negligible but, for the linearly increasing number of changepoints scenario, the difference is relatively large.

A final way to compare the accuracy of PELT with that of BS is to look at how accurately each method detects the actual times at which changepoints occurred. For

Figure 3: Proportion of correctly identified changepoints against the proportion of falsely detected changepoints. Change in variance with  $m = n/100$  where (a)  $n = 500$ , (b)  $n = 5,000$ , (c)  $n = 50,000$  (PELT: thick line, BS: thin line, +: SIC penalty).



the purposes of this study a changepoint is considered correctly identified if we infer its location within a distance of 10 time-points of the true position. If two changepoints are identified in this window then one is counted as correct and one false. The number of false changepoints is then the total number of changepoints identified minus the number correctly identified. The results are depicted in Figure 3 for a selection of data lengths,  $n$ , for the case  $m = n/100$ . As  $n$  increases the difference between the PELT and BS algorithms becomes clearer with PELT correctly identifying more changepoints than BS. Qualitatively similar results are obtained if we change how close an inferred changepoint has to be to a true changepoint to be classified as correct. Figures for square root increasing and fixed numbers of changepoints are given in the supplementary material. As the number of changepoints decreases a higher proportion of true changepoints are detected with fewer false changepoints. The supplementary material also contains an exploration of the same properties for changes in both mean and variance. The results are broadly similar to those described above. We now demonstrate increased accuracy of the PELT algorithm compared with BS on an oceanographic data set; a financial application is given in the supplementary material.

## 4.2 Application to Canadian Wave Heights

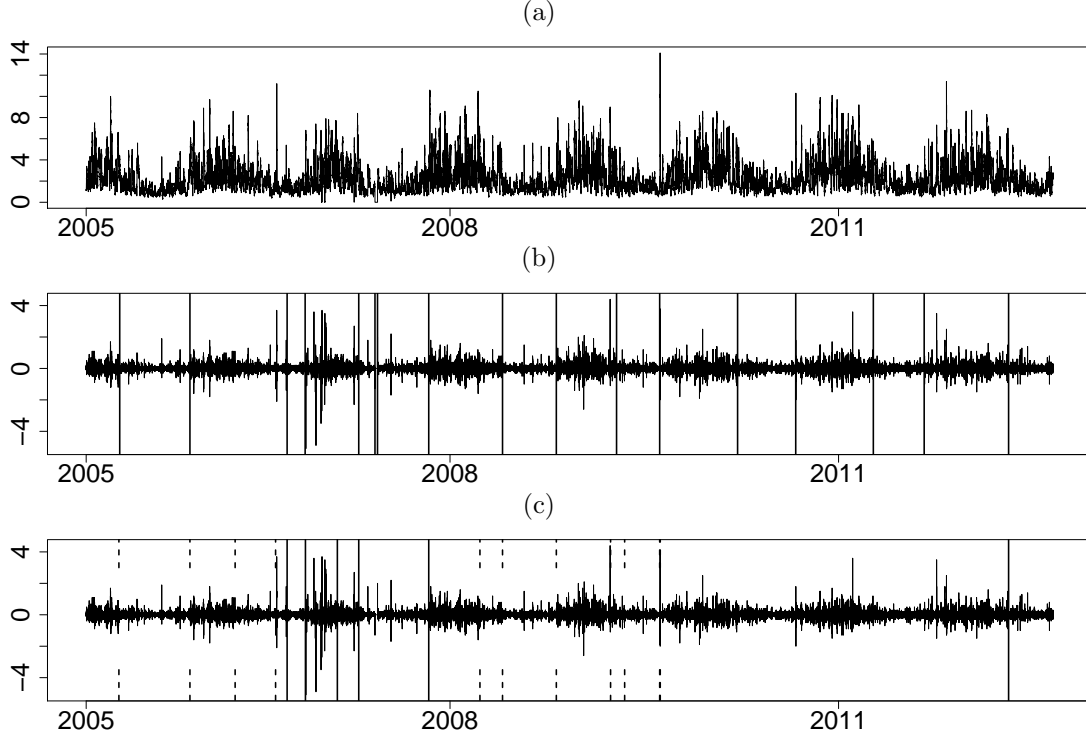
There is interest in characterising the ocean environment, particularly in areas where there are marine structures, e.g. offshore wind farms or oil installations. Short-term operations, such as inspection and maintenance of these marine structures, are typically performed in periods where the sea is less volatile to minimize risk.

Here we consider publically available data for a location in the North Atlantic where data has been collected on wave heights at hourly intervals from January 2005 until September 2012, see Figure 4(a). Our interest is in segmenting the series into period of lower and higher volatility. The data we use is obtained from Fisheries and Oceans Canada, East Scotian Slope buoy ID C44137 and has been reproduced in the `changepoint` R package (Killick and Eckley, 2010).

The cyclic nature of larger wave heights in the winter and small wave heights in the summer is clear. However, the transition point from periods of higher volatility (winter storms) to lower volatility (summer calm) is unclear, particularly in some years. To identify these features we work with the first difference data. Consequently a natural approach is to use the change in variance cost function of Section 4.1. Of course this is but one of several ways in which the data could be segmented.

For the data we consider (Figure 4(a)) there is quite a difference in the number of changepoints identified by PELT (17) and optimal Binary Segmentation (6). However, the location of the detected changepoints is quite similar. The difference in likelihood between the inferred segmentations is 3851. PELT chooses a segmentation which, by-eye, segments the series well into the different volatility regions (Figure 4(b)). Conversely, the segmentation produced by BS does not (Figure 4(c)); most notably it fails to detect any transitions between 2008 and 2012. If we increase the number of changepoints BS finds to equal that of PELT, the additional changepoints still fail to capture the regions appropriately.

Figure 4: North Atlantic Wave Heights (a) Original data (b) Differenced data with PELT change-points (c) Differenced data with optimal BS changepoints and additional subBS changepoints (dotted lines).



#### 4.3 Changes in Auto-covariance within Autoregressive Data

Changes in AR models have been considered by many authors including Davis et al. (2006), Huskova et al. (2007) and Gombay (2008). This section describes a simulation study that compares the properties of PELT and the genetic algorithm used in Davis et al. (2006) to implement the minimum description length (MDL) test statistic.

***Minimum Description Length for AR Models*** The simulation study here will be constructed in a similar way to that of Section 4.1. It is assumed that the data follow an autoregressive model with order and parameter values depending on the segment. We shall take the cost function to be the MDL, and consider allowing AR models of order  $1, \dots, p_{max}$ , for some chosen  $p_{max}$  within each segment. The

associated cost for a segment is

$$\mathcal{C}(y_{(\tau_{i-1}+1):\tau_i}) = \min_{p \in \{1, \dots, p_{max}\}} \left\{ \log p + \frac{p+2}{2} \log(\tau_i - \tau_{i-1}) + \frac{\tau_i - \tau_{i-1}}{2} \log(2\pi\hat{\sigma}(p, \tau_{i-1} + 1, \tau_i)^2) \right\}. \quad (11)$$

where  $\hat{\sigma}(p, \tau_{i-1} + 1, \tau_i)^2$  is the Yule-Walker estimate of the innovation variance for data  $y_{(\tau_{i-1}+1):\tau_i}$  and order  $p$ . When implementing PELT, we set  $K = -[2 \log(p_{max}) + (p_{max}/2) \log(n)]$ , to ensure that (4) is satisfied.

**Simulation Study** The simulated data consists of 5 scenarios with varying lengths,  $n = c(1000, 2000, 5000, 10000, 20000)$  and each scenario contains  $0.003n$  change-points. These changepoints are distributed uniformly across  $(2, n - 2)$  with the constraint that there must be at least 50 observations between changepoints. Within each of these 5 scenarios we have 200 repetitions where the segment order is selected randomly from  $\{0, 1, 2, 3\}$  and the autoregressive parameters for each segment are a realisation from a standard Normal distribution subject to stationarity conditions. We compare the output from PELT with an approximate method proposed by Davis et al. (2006) for minimising the MDL criteria, which uses a genetic algorithm. This was implemented in the program Auto-PARM, made available by the authors. We used the recommended settings except that for both methods we assumed  $p_{max} = 7$ . Table 1 shows the average difference in MDL over each scenario for each fitted model. It is clear that on average PELT achieves a lower MDL than the Auto-PARM algorithm and that this difference increases as the length of the data increases. Overall, for 91% of data sets, PELT gave a lower value of MDL than Auto-Parm. In addition, the average number of iterations required for PELT to converge is small in all cases. Previously, it was noted that the PELT algorithm for the MDL penalty is no longer an exact search algorithm. For  $n = 1,000$  we evaluated the accuracy of PELT by calculating the optimal segmentation in each case using Segment Neighbourhood

Table 1: Average MDL and number of PELT iterations over 200 repetitions.

$n$	1,000	2,000	5,000	10,000	20,000
no. iterations	2.470	2.710	2.885	2.970	3.000
Auto-PARM - PELT	8.856	13.918	59.825	252.796	900.869

(SN). The average difference in MDL between the SN and PELT algorithms is 1.01 (to 2dp). However SN took an order of magnitude longer to run than PELT, its computational cost increasing with the cube of the data size making it impracticable for large  $n$ . A better approach to improve on the results of our analysis would be to improve the search strategy for the value of penalty function to run PELT with.

## 5. DISCUSSION

In this paper we have presented the PELT method; an alternative exact multiple changepoint method that is both computationally efficient and versatile in its application. It has been shown that under certain conditions, most importantly that the number of changepoints is increasing linearly with  $n$ , the computational efficiency of PELT is  $\mathcal{O}(n)$ . The simulation study and real data examples demonstrate that the assumptions and conditions are not restrictive and a wide class of cost functions can be implemented. The empirical results show a resulting computational cost for PELT that can be orders of magnitude smaller than alternative exact search methods. Furthermore, the results show substantial increases in accuracy by using PELT compared with Binary Segmentation. Whilst PELT is not, in practice, computationally quicker than Binary Segmentation, we would argue that the statistical benefits of an exact segmentation outweigh the relatively small computational costs. There are other fast algorithms for segmenting data that improve upon Binary Segmentation (Gey and Lebarbier, 2008; Harchaoui and Levy-Leduc, 2010), although these do not have the guarantee of exactness that PELT does.

Rigaiil (2010) develops a competing exact method called pruned dynamic programming (PDPA). This method also aims to improve the computational efficiency of an exact method, this time Segment Neighbourhood, through pruning, but the way pruning is implemented is very different from PELT. The methods are complementary. Firstly they can be applied to different problems, with PDPA able to cope with a non-linear penalty functions for the number of changepoints, but restricted to models with a single parameter within each segment. Secondly the applications under which they are computationally efficient is different, with PDPA best suited to applications with few changepoints. Whilst unable to compare PELT with PDPA on the change in variance or the change in mean and variance models considered in the results section, we have done a comparison between them on a change in mean. Results are presented in Table 1 of the Supplementary material. Our comparison was for both a linearly increasing number of changepoints, and a fixed number of changepoints scenario. For the former PELT was substantially quicker, by a factors of between 300 and 40,000 as the number of data-points varied between 500 and 500,000. When we fixed the number of changepoints to 2, PDPA was a factor of 2 quicker for data with 500,000 changepoints, though often much slower for smaller data sets.

Code implementing PELT is contained within the R library `changepoint` which is available on CRAN (Killick and Eckley, 2010).

## References

- Aggarwal, R., Inflan, C., and Leal, R. (1999). Volatility in emerging stock markets. *The Journal of Financial and Quantitative Analysis*, 34(1):33–55.
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.

- Andreou, E. and Ghysels, E. (2002). Detecting multiple breaks in financial market volatility dynamics. *Journal of Applied Econometrics*, 17(5).
- Auger, I. E. and Lawrence, C. E. (1989). Algorithms for the optimal identification of segment neighborhoods. *Bulletin of Mathematical Biology*, 51(1):39–54.
- Bellman, E. and Dreyfus, S. E. (1962). *Applied Dynamic Programming*. Princeton.
- Birge, L. and Massart, P. (2007). Minimal Penalties for Gaussian Model Selection. *Probability Theory and Related Fields*, 138(1):33–73.
- Braun, J. V., Braun, R. K., and Muller, H. G. (2000). Multiple changepoint fitting via quasilikelihood, with application to DNA sequence segmentation. *Biometrika*, 87(2):301–314.
- Chen, J. and Gupta, A. K. (2000). *Parametric statistical change point analysis*. Birkhauser.
- Davis, R. A., Lee, T. C. M., and Rodriguez-Yam, G. A. (2006). Structural break estimation for nonstationary time series models. *JASA*, 101(473):223–239.
- Eckley, I. A., Fearnhead, P., and Killick, R. (2011). Analysis of changepoint models. In Barber, D., Cemgil, T., and Chiappa, S., editors, *Bayesian Time Series Models*. Cambridge University Press.
- Fernandez, V. (2004). Detection of breakpoints in volatility. *Estudios de Administracion*, 11(1):1–38.
- Gey, S. and Lebarbier, E. (2008). Using cart to detect multiple change-points in the mean for large samples. *SSB Preprint, No.12*.
- Gombay, E. (2008). Change detection in autoregressive time series. *J. Multivariate Anal.*, 99(3):451 – 464.

- Guyon, X. and Yao, J.-f. (1999). On the underfitting and overfitting sets of models chosen by order selection criteria. *J. Multivariate Anal.*, 70(2):221 – 249.
- Harchaoui, Z. and Levy-Leduc, C. (2010). Multiple change-point estimation with a total-variation penalty. *JASA*, 105(492):1480–1493.
- Horvath, L. (1993). The maximum likelihood method of testing changes in the parameters of normal observations. *Annals of Statistics*, 21(2):671–680.
- Huskova, M., Praskova, Z., and Steinebach, J. (2007). On the detection of changes in autoregressive time series. *J. Statist. Plann. Inference*, 137(4):1243 – 1259.
- Inclan, C. and Tiao, G. C. (1994). Use of cumulative sums of squares for retrospective detection of changes of variance. *JASA*, 89(427):913–923.
- Jackson, B., Sargle, J. D., Barnes, D., Arabhi, S., Alt, A., Gioumoussis, P., Gwin, E., Sangtrakulcharoen, P., Tan, L., and Tsai, T. T. (2005). An algorithm for optimal partitioning of data on an interval. *IEEE, Signal Processing Letters*, 12(2):105–108.
- Killick, R. and Eckley, I. A. (2010). *changepoint: Analysis of Changepoint Models*. Lancaster University, Lancaster, UK.
- Olshen, A. B., Venkatraman, E. S., Lucito, R., and Wigler, M. (2004). Circular binary segmentation for the analysis of array-based DNA copy number data. *. Biostatistics*, 5:557–72.
- Picard, F., Lebarbier, E., Hoebeke, M., Rigai, G., Thiam, B., and Robin, S. (2011). Joint segmentation, calling and normalization of multiple cgh profiles. *Biostatistics*, 12(3):413 – 428.
- Picard, F., Robin, S., Lavielle, M., Vaisse, C., and Daudin, J. J. (2005). A statistical approach for array cgh data analysis. *Bioinformatics*, 6(27).

- Rigaiil, G. (2010). Pruned dynamic programming for optimal multiple change-point detection. *ArXiv e-prints*.
- Schwarz, G. (1978). Estimating the dimension of a model. *Ann. Statist.*, 6(2):461–464.
- Scott, A. J. and Knott, M. (1974). A cluster analysis method for grouping means in the analysis of variance. *Biometrics*, 30(3):507–512.
- Sen, A. and Srivastava, M. S. (1975). On tests for detecting change in mean. *The Annals of Statistics*, 3(1):98–108.
- Yao, Y. (1984). Estimation of a noisy discrete-time step function: Bayes and empirical Bayes approaches. *The Annals of Statistics*, 12(4):1434–1447.
- Zhang, N. R. and Siegmund, D. O. (2007). A modified bayes information criterion with applications to the analysis of comparative genomic hybridization data. *Biometrics*, 63(1):22–32.